

The Emergence of Organizing Structure in Conceptual Representation

Brenden M. Lake,¹ Neil D. Lawrence,² Joshua B. Tenenbaum,^{3,4}

¹Center for Data Science, New York University

²Department of Computer Science, University of Sheffield

³Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

⁴Center for Brains Minds and Machines

Abstract

Both scientists and children make important structural discoveries, yet their computational underpinnings are not well understood. Structure discovery has previously been formalized as probabilistic inference about the right structural form — where form could be a tree, ring, chain, grid, etc. [Kemp & Tenenbaum (2008)]. The discovery of structural form. *PNAS*, 105(3), 10687-10692]. While this approach can learn intuitive organizations, including a tree for animals and a ring for the color circle, it assumes a strong inductive bias that considers only these particular forms, and each form is explicitly provided as initial knowledge. Here we introduce a new computational model of how organizing structure can be discovered, utilizing a broad hypothesis space with a preference for sparse connectivity. Given that the inductive bias is more general, the model’s initial knowledge shows little qualitative resemblance to some of the discoveries it supports. As a consequence, the model can also learn complex structures for domains that lack intuitive description, as well as predict human property induction judgments without explicit structural forms. By allowing form to emerge from sparsity, our approach clarifies how both the richness and flexibility of human conceptual organization can coexist.

1 Introduction

Structural discoveries play an important role in science and cognitive development (Carey, 2009; Kuhn, 1962). In biology, Linnaeus realized that living things were best organized as a tree,

displacing the “great chain of being” used for centuries before. Modern chemistry began with Mendeleev’s discovery of the periodic structure of elements. In cognitive development, children realize that the days of the week form a cycle and that social networks form cliques. Children do not initially treat comparative relations such as “longer than” as transitive, but they realize the unidimensional structure around age 7 or 8 (Inhelder & Piaget, 1964). While learning the names for objects, children tend to learn just one name per object, but they must later discover that names can be organized in a taxonomy (Markman, 1989). Yet, other forms of learning are not marked by clear structural insights and transitions. Everyday forms of learning, such as a child acquiring facts about animals and artifacts, have been characterized as accumulating correlations and gradually differentiating concepts (Rogers & McClelland, 2004). What distinguishes more ordinary forms of learning from genuine structural discoveries?

This distinction is not well understood computationally. In machine learning, many standard algorithms learn only a single form of structure that is decided on in advance – whether it is learning clusters (e.g., k-means), trees (hierarchical clustering), or multidimensional spaces (principal component analysis). It is usually up to the human practitioner, rather than the algorithms themselves, to decide which type of structure is most appropriate for the data. Other approaches such as neural networks (deep learning; LeCun, Bengio, & Hinton, 2015) make very general representational assumptions and can model data with many different structural characteristics. However, the learned structure is only implicit in the connection weights, with no obvious mechanism for differentiating between genuine structural discoveries from other forms of learning.

Previous work from our group investigated how the form of structure can be learned from data and represented explicitly (Kemp & Tenenbaum, 2008). Structure selection can be formalized as probabilistic inference over a set of discrete and mutually exclusive hypotheses such as rings (i), grids (ii), trees (iii), chains (iv), etc. called *structural forms* (Fig. 1A). The forms are defined by grammatical constraints on the connections between entities; for example the ring form constrains each node to have exactly two neighbors. This approach can learn intuitive structures for a variety of domains, such as a tree for animals and a ring for the color circle. Once learned, these structures can be used to explain human inductive reasoning about novel properties of objects (Kemp & Tenenbaum, 2009).

While powerful, the structural forms approach is not appropriate when the data strays from the

predefined forms. Exceptions are common in real world domains. While the genetic similarity of animals is captured by an evolutionary tree,¹ everyday reasoning about animals draws on factors that span divergent branches, including shared habitat, role as predator versus prey, and size. While these factors cannot be perfectly explained by a single tree, other domains are interestingly structured yet even further removed from a pristine form, such as artifacts and social networks. Since people can learn and reason in all of these domains, they must either entertain structural hypotheses without strict grammatical constraints, or engage in other types of learning.

An alternative approach is to learn implicit rather than explicit structural organization. Rogers and McClelland (2004) studied a neural network that learns to map animals (like canary) and relations (can) to output attributes that a canary can do (grow, move, fly, and sing). Like the structural forms approach, the neural network learns aggregate statistical structure from the observations. Rogers and McClelland (2004) analyzed their network’s learned representation through dimensionality reduction, projecting each living thing into a low-dimensional representational space. In this space, the most salient split is between animals versus plants, with sub-divisions for mammals versus birds, trees versus flowers, etc. While the network is not constrained or biased to learn a tree, it nonetheless learns a distributed representation for living things that exhibits hierarchical structure.

While neural networks are powerful implicit learners, people also seem to learn and reason with explicit structural representations. Kemp and Tenenbaum (2009) points out that language often carries direct structural information such as “Indiana is next to Illinois” or “A dolphin is a mammal not a fish,” observations that can be elegantly incorporated in more explicit representations. These representations can also provide scaffolding for learning higher-level concepts with direct structural interpretations, such as using a tree structure to help learn the word “primate” (Xu & Tenenbaum, 2007). Moreover, it remains mysterious how implicit structure might, in rare yet pivotal moments in science and childhood, crystallize to create explicit structural discoveries.

Here, we present a new computational approach to structure learning and discovery that incorporates some of the best features of previous probabilistic and connectionist models. Rather than selecting between discrete and mutually exclusive structural hypotheses defined by grammars, the

¹Even this structure has exceptions; for example, Rivera and Lake (2004) provide evidence that at the deepest levels “the tree of life is actually a ring of life” where genomes fused.

model learns explicit structure in an unrestricted hypothesis space of all possible graphs. A key insight is that many cognitively natural structural forms are very sparse, meaning that the graphs have relatively few edges (connecting lines between the nodes). This observation is central to our new approach, called *structural sparsity*, which is capable of rich structural inferences guided by a preference for sparsity. This not only suggests an alternative means for learning pristine structural form (Fig. 1A), but it also explains how learning could proceed when some domains, such as artifacts, stray from the cognitively natural forms (Fig. 1B).

We also use the comparison between structural sparsity and structural forms as an opportunity to explore two different types of inductive bias and their roles in supporting learning. For any learning algorithm, it is instructive to compare the initial knowledge (before data) with the final knowledge (after data), as means of distinguishing what the model learns from what the model starts with. For the structural forms model (Fig. 2 left), samples from the prior are already highly structured graphs (rings, grids, trees, etc.), suggesting that the primary role of learning is to select one of the provided forms, while simultaneously organizing the entities subject to the corresponding grammatical constraints. In contrast, for the structural sparsity model (Fig. 2 right), the model’s initial knowledge shows little qualitative resemblance to many of the structures it discovers, resembling some of the representational leaps in science and cognitive development (Carey, 2009). By comparing structural forms and structural sparsity on the same data sets – two closely related models with different inductive biases – we can distinguish between the cases where a strong prior over forms is needed, from other cases where a more general bias towards sparsity can lead to rich structural insight.

2 Model

2.1 Structural sparsity

Bayesian learning underlies both the structural forms and structural sparsity models, which use Bayes’ rule to reason about the distribution of structural hypotheses S in light of data D ,

$$P(S|D) \propto P(D|S)P(S). \quad (1)$$




Figure 1: Structures that can be represented (A; i: ring, ii: grid, iii: tree, iv: chain), or cannot be represented (B), by the structural forms of Kemp and Tenenbaum (Kemp & Tenenbaum, 2008). All of the structures (A & B) are sparse, meaning they have relatively few edges (lines) between the nodes. Sparsity is the foundation of the computational model introduced in this work.

Both models use a similar likelihood $P(D|S)$ with different priors $P(S)$. The prior for the structural forms model is a mixture of discrete forms F (such as rings, trees, grids, etc.), where $P(S) = \sum_F P(S|F)P(F)$. The initial knowledge $P(S)$ reflects the fact that the forms are provided explicitly and no other possibilities are entertained (Fig. 2 left). The prior for structural sparsity covers a broader hypothesis space that includes all possible connectivity configurations (Fig. 2 right). Sparsity is operationalized with an improper prior that decreases the score with each additional edge in a structure

$$P(S) \propto \exp(-\beta(\#S)), \quad (2)$$

where $\#S$ is the number of edges in a structure. In both models, S specifies connections between clusters of objects. Each object connects to exactly one latent cluster, called a cluster node (Kemp & Tenenbaum, 2008) (gray nodes in Fig. 2), and cluster nodes are non-empty. Using cluster nodes, groups of similar objects can be elegantly represented without dense object to object connections, allowing for sparser graphs that highlight the underlying topology. Besides the cluster assignments, objects do not form any other direct connections, unlike an earlier version of our model without cluster nodes (Lake & Tenenbaum, 2010). Given that the number of edges typically grows with the number of clusters, the structural sparsity prior tends to favor both fewer edges and fewer clusters.

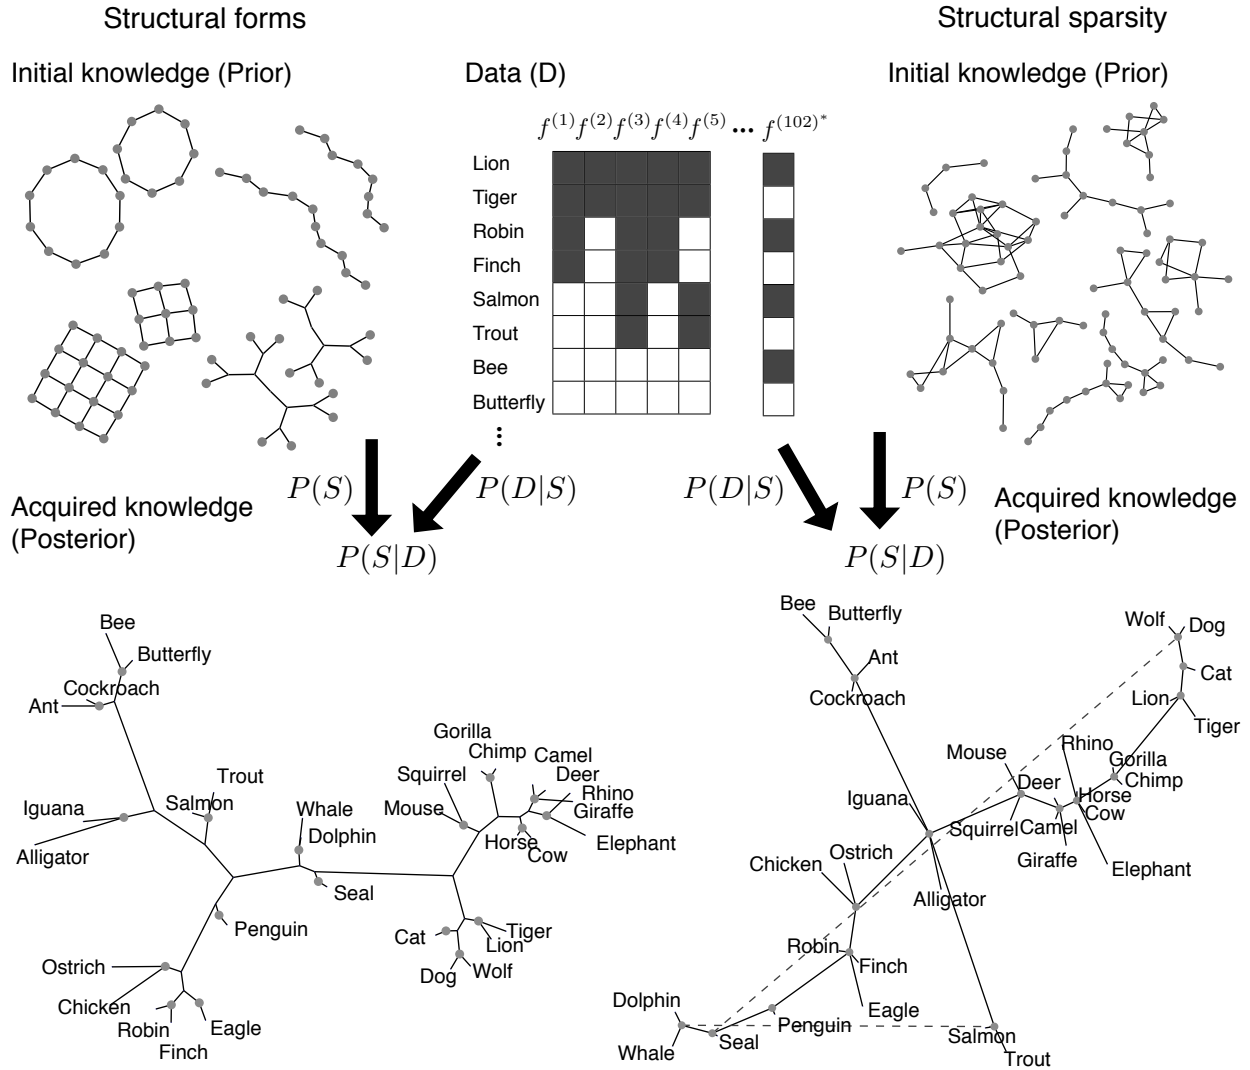


Figure 2: Two models of structure discovery: (left) structural forms (Kemp & Tenenbaum, 2008) and (right) structural sparsity (this work). Bayesian learning of structure from data D involves an update from prior $P(S)$ to posterior $P(S|D)$. Given a matrix of animals and their features, the two models learned similar structures despite the substantial difference in initial knowledge (prior). In the visualizations, stronger edges are shorter, except for the dotted lines that are edges that lie outside the strongest spanning tree.

The term $P(D|S)$ specifies the link between the graph structure and the features of objects. The data $D = \{f^{(1)}, \dots, f^{(m)}\}$ is a matrix where rows are objects and columns are features $f^{(i)}$. For instance, we used a data set of 33 animals and 102 features, where the first five features were “has lungs,” “has a large brain,” “has a spinal cord,” “is warm-blooded,” and “has teeth” (Fig. 2 top). The distribution $P(D|S) = \prod_{i=1}^m P(f^{(i)}|S)$ treats features as independent draws from a distribution that prefers smoothness, meaning objects tend to share similar feature values with their neighbors in the graph. For example, all the features shown in Fig. 2 are relatively smooth (except for a hypothetical $f^{(102*)}$ which is highly non-smooth) over the graphs that represent animals (Fig. 2 bottom). An edge between node i and j is parameterized by a continuous value $s_{ij} > 0$ that determines its strength, and absent edges have $s_{ij} = 0$. The graph as a whole specifies the covariance matrix of a multivariate Gaussian distribution with one dimension per node (Appendix 8.1). While $P(D|S)$ is best suited for continuous features, binary data can be treated as continuous (Kemp & Tenenbaum, 2008). Data in the form of a similarity matrix can also be accommodated by treating it as a covariance matrix and then sampling an input data matrix D (we used 2000 features).

As in the structural forms model, we are interested in the single best structure given the data. For structural sparsity, taking the log of the posterior probability (Eq. 1) leads to the optimization problem

$$\operatorname{argmax}_S \sum_{i=1}^m \log P(f^{(i)}|S) - \beta(\#S), \quad (3)$$

which emphasizes the trade-off between capturing regularities in the features (left term) while promoting the sparsity of the structure (right term). There is only one free parameter β which controls the amount of sparsity.

2.2 Search algorithm

Searching for the maximum of this score (Eq. 3) is a special case of learning the structure of an undirected graphical model, a generally intractable problem that requires heuristic algorithms not guaranteed to reach the global optimum (Koller & Friedman, 2009). Here we describe the heuristic search algorithm developed for solving this problem. At its core, structural sparsity is a proposal for the computational level problem (Marr, 1982) of structure discovery, and thus the particular

search algorithm we used was not intended to be cognitive plausible.

Our search algorithm consists of an outer-routine that looks for the best partition of objects into cluster nodes, as well as a sub-routine that evaluates a candidate partition by optimizing for the best set of edges. The outer-routine tries local proposals for splitting and merging clusters and picks the best move at each step (Kemp & Tenenbaum, 2008). While searching over partitions can be challenging, it can be sped up greatly by a preprocessing stage that clusters objects several times, with different numbers of clusters, and picks the best scoring option for initialization.

The sub-routine has the challenging task of estimating the sparse connectivity between clusters, necessary for evaluating partition quality during both initialization and local search. Directly searching the combinatorial space of sparsity patterns is intractable, but there are well-known heuristics such as ℓ_1 that push parameters to zero ($s_{ij} = 0$) and consequently propose a sparsity pattern (Tibshirani, 1996). The ℓ_1 relaxation penalizes the sum rather than the number of parameters $\#S \approx \lambda \sum_{ij} |s_{ij}|$ for some constant λ . Learning a Gaussian graph with ℓ_1 penalization from a complete data matrix is a convex optimization problem that can be solved efficiently (Banerjee, El Ghaoui, & D’Aspremont, 2008; Lake & Tenenbaum, 2010). Although our connectivity estimation problem is not convex due to the latent nodes and missing data, efficient solvers for the complete data case are utilized in a Structural Expectation-Maximization algorithm guaranteed to improve the primary objective in Eq. 3 with each iteration (Friedman, 1997).

The Appendix (Section 8) contains a more detailed description of the entire search algorithm. Although search can be slow, it considers the full range structures, unlike the implementation of structural forms that runs a separate search for each form (Kemp & Tenenbaum, 2008). Code for running the model is available online.²

3 Discovering structure in synthetic data

The algorithm was evaluated on its ability to recover the true topology of synthetic data. Given that objects can cluster in a variety of patterns atop these forms, the challenge of jointly finding the right cluster partition and the right sparsity pattern results in a very difficult combinatorial search problem. Structural sparsity was tested on both its ability to recover structural forms (Fig.

²<https://github.com/brendenlake/structural-sparsity>




Figure 3: Synthetic data sets were generated from structures with one object per cluster node (A) and arbitrary clusters of objects on the underlying structure (B).

1A) as well as novel structures that share interesting features with the forms, but not the pristine qualities necessary for identification with the Kemp and Tenenbaum (2008) approach (Fig. 1B). These new structures include a tree embedded in a ring (i), a non-grid planar graph (ii), a ring of trees (iii), and multiple disconnected forms including disjoint chains (iv; Fig. 1B).

The synthetic experiments were structured as follows. For each candidate synthetic structure, 1000 features were generated from the ground truth graph and provided to the model. The sparsity free parameter was set to $\beta = 6$. Given that the algorithm is non-deterministic, the search process was repeated 10 times for each synthetic data set. Each run of the search algorithm is compared to the ground truth.

The structures in the first set of experiments had exactly one object per cluster node (Fig. 3A). This clustering pattern is easier for the search algorithm to find compared to an arbitrary clustering, in part because it is automatically examined as a candidate pattern for initialization. Thus, this helps to isolate the challenge of finding the right sparsity pattern from the problem of finding the right clustering. The results of search are shown Table 1. In each case, algorithm

recovers a structure that scores (Eq. 3) at least as well as the original structure (Table 1 Column 1, score match).³ For all cases except the chain, the highest scoring structure recovered the right partition of objects into cluster nodes (Table 1 Column 2, partition match), and furthermore the structure matched the adjacency matrix up to a permutation of the cluster node indices (Table 1 Column 3, exact match). There is straightforward reason why the chain structure was not recovered exactly. To increase sparsity, the model tended to combine the last two objects on each end of the chain into a single cluster, using a weaker connection to the object at the end of the chain. This configuration retains the correct relations while increasing sparsity, and thus the learned structure scores higher than the original graph.⁴

The structures in the second set of experiments had richer clustering patterns on the same graphs, as well as others like clusters, a tree, a peace sign, disjoint chains, a ring of trees, and a plane (Fig. 3B). The results are shown in Table 2. The clusters, ring, tree, and ring of trees were perfectly recovered for each run. Like in the singleton case, the algorithm always conserved edges on the ends of the chain structure and disjoint chains structure. Mistakes occurred on 1 run of the grid, 2 runs of the peace graph, 1 run of the plane, and all of these were unnecessary cluster node splits that did not alter the underlying topology. For several runs of the plane structure, the model recovered an alternative graph that differed in the location of one edge, and it also scored roughly as well as the ground truth. On the whole, our search algorithm was highly successful at recovering ground truth structure from synthetic data.

	Score match	Partition match	Exact match
Ring	10	10	10
Chain	10	0	0
Grid	10	10	10
Peace	10	10	10

Table 1: Synthetic experiments with singleton cluster nodes (Fig. 3A). Entries show the number of successes out of 10 runs of the search algorithm.

³Structures are counted as correct if they are within 4 log points of the ground truth, so that an extra edge will be counted as an error (-6 log points).

⁴This reduction also occurred in the original structural forms model, in order to use fewer cluster nodes.

	Score match	Partition match	Exact match
Clusters	10	10	10
Ring	10	10	10
Chain	10	0	0
Tree	10	10	10
Grid	9	9	9
Peace	8	8	8
Disjoint	10	0	0
chains			
Ring	10	10	10
of trees			
Plane	9	9	3

Table 2: Synthetic experiments with multi-object cluster nodes (Fig. 3B). Entries show the number of successes out of 10 runs of the search algorithm.

4 Discovering structure in empirical data

Structural sparsity was applied to six empirical data sets. To facilitate comparison with the structural forms approach, four data sets from Kemp and Tenenbaum (2008) were included. The search algorithm was run ten times for each data set and the highest scoring structure is displayed. The sparsity prior was fixed at $\beta = 6$. In each case, the output of the structural forms model is displayed for comparison. The data sets were rescaled following the procedure of Kemp and Tenenbaum (2008).⁵

4.1 Animals

The first data set is shown in Fig. 2 and contains a set of 33 animals with 102 features related to biology, anatomy, and habitat (collected by Kemp & Tenenbaum, 2008). Biological features included those like “has lungs” and “is warm-blooded”, anatomical features included “has fins” and “has a long neck,” and perceptual features included “is black.” Habitat features often highlight non-biological similarity, like the fact that fish and the aquatic mammals “live in the ocean.”

While structural forms learned a tree, structural sparsity learned a graph with a tree-based backbone that primarily reflects biology, with branches for the insects, fish, birds, and mammals. Structural sparsity can also capture relationships that reach across the tree. The aquatic mammals

⁵Each data matrix D (n_x objects by m features) was linearly transformed so the mean across the entire matrix is 0 and the largest element in $\frac{1}{m}DD^T$ is 1. If there are missing values, the features with the most common pattern of missing objects are grouped, and the transform is computed for this sub-matrix and applied to the full matrix.

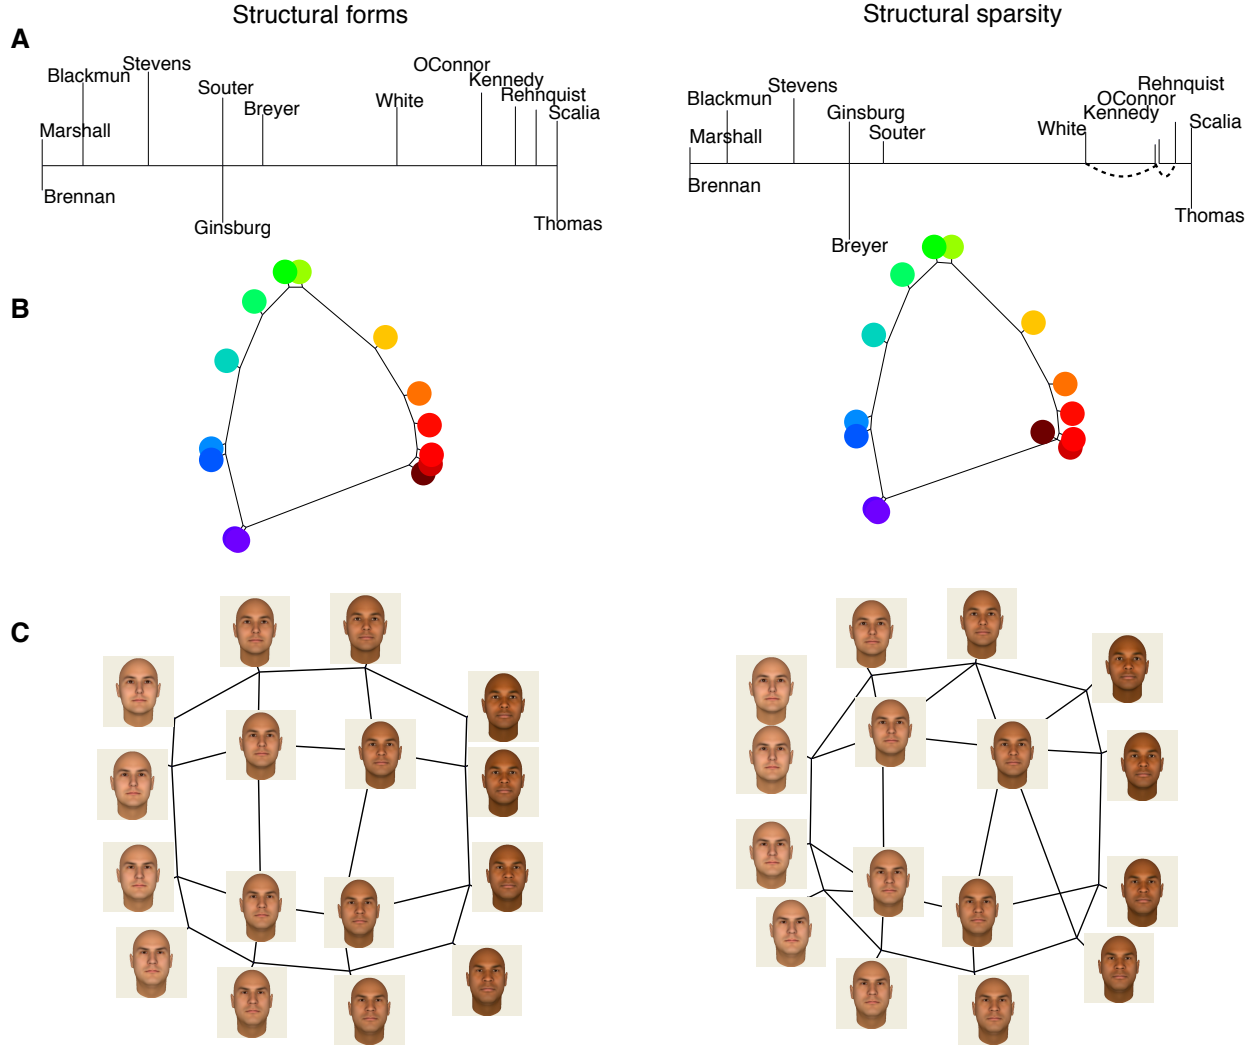


Figure 4: Structures learned by structural forms (left) and sparsity (right) for a vote matrix from 13 Supreme Court justices (A), similarity judgments between spectral colors (B), and Euclidean distance between images of faces (C). To visualize the spatial structure, node locations were chosen to be the principal eigenvectors of the implied covariance matrices (see Lawrence, 2011, 2012). Only one dimension is shown in (A), and dotted lines are edges that do not follow a rigid linear order.

are situated with the fish, penguin, and other mammals, representing both biology, shared habitat, and visual similarity. The tree-based structural form better reflects evolutionary branching, where objects must lie at the leaf nodes. Trees such as this are not considered by structural sparsity, since allowing empty cluster nodes would make search much more difficult. But these form-based latent trees are not necessarily highly sparse, and here structural sparsity found about half as many non-attachment edges.

4.2 Judges

The second data set included 13 Supreme Court justices and their votes on 1,596 cases (Fig. 4A), collected and preprocessed by Kemp and Tenenbaum (2008) to include 1596 cases under Chief Justice Rehnquist. Due to non-participation and the fact that at most 9 justices serve at a time, there were many missing values. We integrated over the missing values while computing the model score.

Both models organized the justices along a spectrum from liberal (Marshall and Brennan) to conservative (Scalia and Thomas). While structural forms is constrained to learn exactly a chain, structural sparsity learned a similar chain with a few additional relationships between the conservatives (two additional dotted edges and one gap).

4.3 Colors

The next data set includes perceived similarity ratings between 14 spectral colors. Originally collected by (Ekman, 1954), the similarity matrix was published in Shepard (1980) and used in Kemp and Tenenbaum (2008), while assuming a value of 1 along the diagonal. Using this data, the structural sparsity model discovered the color circle first described by Newton (Fig. 4B). While the ring topology was pre-specified in the forms model, it is an emergent consequence of the data combined with structural sparsity.

4.4 Faces

Another data set is based on pixel similarity between images of faces that vary along a masculinity and a race dimension. This similarity data set was created by Kemp and Tenenbaum (2008). The faces vary in two dimensions, race and gender, with four values along each. The similarity

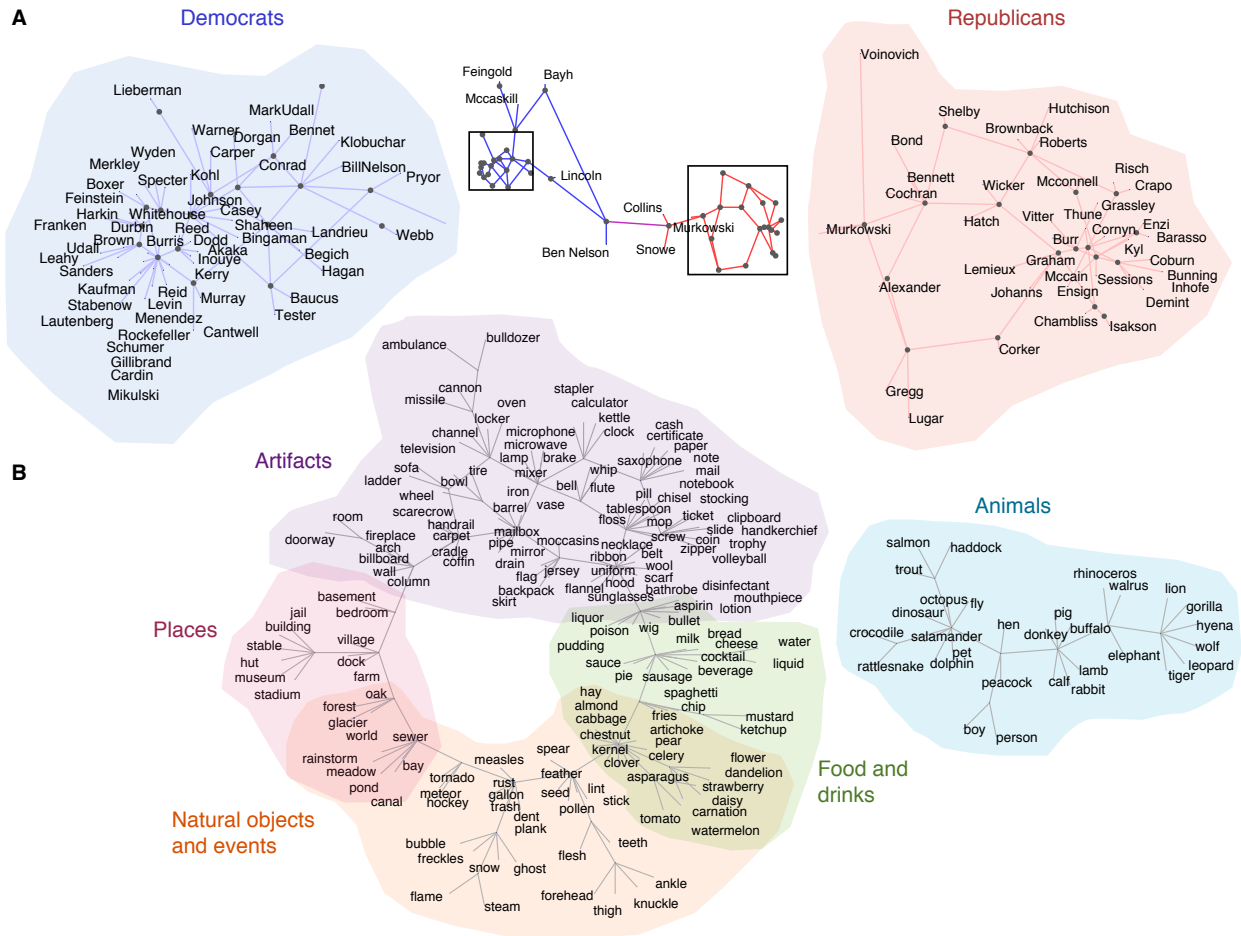


Figure 5: Sparse structures learned from two data sets: US Senators and their votes in the 111th congress (A), and 200 objects and their features with broad semantic coverage (B). In (A), shorter edges are stronger, but in (B) all edges are shown at approximately the same length for visual clarity. Some coherent sub-regions are shaded and labeled for easier interpretation; these are not labels provided to the unsupervised learning algorithm.

matrix between faces is based on the Euclidean distance in pixel space, with 1 along the diagonal. Structural sparsity recovered these two dimensions, although it deviates from a strict grid (Fig. 4C).

4.5 Senate

Structural sparsity can also be applied to domains that have a rich organization but do not fit a cognitively natural structural form. We used a data set of Senators and their voting records was for the 111th United States Congress, from January 2009 to January 2011.⁶ Voting present or not

⁶The votes were retrieved from <http://www.voteview.com/house111.htm>. The various yes votes (1,2,3) and no votes (4,5,6) were collapsed to yes (1) or no (0).

voting was treated as missing data. Only senators that were present for at least half the votes were included, resulting in 98 Senators and 696 votes. Due to the size of of this data set (and the following), the sparsity parameter was increased to $\beta = 18$. To maintain tractability, search was also modified to be less thorough, as described in Appendix 8.3.

Given the voting records, the model recovered the central divide between Democrats and Republicans, bridged only by a connection between moderate Democrat Ben Nelson and Republicans Olympia Snowe and Susan Collins (Fig. 5A). While structural forms might capture the party distinction along a line or with clusters, structural sparsity suggests that the complex within-party dynamics are not naturally represented by any one of the pristine forms.

4.6 Common objects

Structural sparsity was also applied to 200 common objects with broad semantic coverage across artifacts and living things (Palatucci, Pomerleau, Hinton, & Mitchell, 2009). Each object was rated according to 218 properties, including questions like “Is it manmade?”, “Can you hold it?”, and “Does it have feet?” Answers were on a 5 point scale from “definitely no” to “definitely yes” conducted on Amazon Turk. The data set was converted to binary form by coding only the most confident response as 1 and the rest as 0. The original data set contained 1000 objects, and a subset of 200 object was randomly selected for tractability.

Given the object and feature data, structural sparsity learns a complex structure with sub-regions for artifacts, animals, food and drinks, natural objects and events, and places (Fig. 5B). Edges instantiate a notion of aggregate semantic relatedness, resembling classical proposals for semantic networks (Collins & Loftus, 1975). Although the data set is prohibitively large for comparison with the structural forms algorithm, the forms would likely miss much of the fine-grained structure.

5 Property induction

Structural sparsity, in addition to revealing organizing form, can be used to predict how people reason about novel properties. First, structures were learned to represent two data sets, one for mammals and one for cities. Second, the learned structures were used to make predictions about

how people reason in a property induction task.

Inductive questions looked like the following: given that a new biological property is true of dolphins, squirrels, and chimpanzees, how likely is it true of all mammals? (Heit, 1998, 2000; Kemp & Tenenbaum, 2009; Osherson, Smith, Wilkie, Lopez, & Shafir, 1990; Rips, 1975). Earlier work using the structural forms approach found an interesting double dissociation (Kemp & Tenenbaum, 2009): people’s ratings of argument strength for mammals were better predicted by a tree than a continuous 2D space, while analogous ratings regarding geographical properties of cities were better predicted by a 2D space than a tree. Here, we investigated whether structural sparsity can capture both patterns of reasoning without requiring special purpose structural forms.

5.1 Data sets for structure learning

The mammals data set consists of 50 mammals and 85 features (properties), primarily collected by Osherson, Stern, Wilkie, Stob, and Smith (1991) and extended in Kemp and Tenenbaum (2009). The features contain biological, anatomical, behavioral, and habitat properties (“is gray”, “has tough skin”, “big”, “swims”, “lives in water”). Features were collected by asking participants to rate the strength of association between each animal and feature, starting at zero with no upper bound. Ratings were scaled between 0 and 100 and then averaged across participants.

The cities data set is a similarity matrix of 9 US cities collected by Kemp and Tenenbaum (2009). Participants were asked to draw the locations of the 9 cities on a piece of paper, ensuring that the relative distances were as accurate as possible. The distances for each participant were scaled so that the largest distance was 1, and then each value was subtracted from 1 to create a measure of spatial similarity. The matrix was averaged across participants and then provided as input to the models.

5.2 Models for structure learning

Four alternative models were trained on the two base data sets: structural sparsity, a tree structural form, a 2D spatial model, and the raw covariance matrix. For each of the models, the product of learning is a multivariate Gaussian distribution with one dimension per object. Each multivariate Gaussian parameterizes a joint distribution on new object properties, which can be used to make Bayesian predictions regarding the strength of inductive arguments. All inductive arguments use

new (generic) properties that were not features in the data set used for structure learning.

- **Structural sparsity.** Structural sparsity represents a Gaussian by a sparse graph (Section 8.1). The structural sparsity model has one free parameter β , which controls the degree of sparsity in the learned structure. Structures were learned across a range of sparsity values. For each value of β , the search algorithm was run twice for mammals and cities, and the highest scoring structure (based on the score in Equation 3, not correlation with participants) was chosen to represent that value of sparsity. Examples of learned structures for mammals are shown in Fig. 7A & D and for cities in Fig. 6B-i.
- **Tree-based structural form.** For both data sets, we used the trees learned by the structural forms model from Kemp and Tenenbaum (2009), with learned structures shown in Fig. 7B and Fig. 6B-ii. Given that the transformation from a graph to a Gaussian differs in some minor ways between structural forms and structural sparsity, we relearned the edge strengths for the tree using maximum likelihood and the graph formalism from structural sparsity. However, re-learning the edge weights did little to change the predictive performance on the property induction tasks.
- **2D spatial form.** We used a spatial model from Kemp and Tenenbaum (2009) to compare with the graph-based methods, with learned representational spaces shown in Fig. 7C and Fig. 6B-iii. This model represents a covariance matrix Σ using distances in a 2D space

$$\Sigma_{ij} = \frac{1}{2\pi} \exp\left(-\frac{1}{\sigma} \|x_i - x_j\|\right), \quad (4)$$

where x_i is the 2D location of object i and $\|x_i - x_j\|$ is the Euclidean distance between two objects.

- **Raw covariance.** The raw covariance model was either identical to the similarity matrix or $\Sigma = \frac{1}{m} DD^T$ where D is the rescaled feature matrix.

5.3 Data sets for property induction

The property induction data concerning mammals, including the Osherson horse and Osherson mammals tasks, was reported in Osherson et al. (1990). Judgments concerned 10 species: horse,

cow, chimp, gorilla, mouse, squirrel, dolphin, seal, and rhino. Participants were shown arguments of the form “Cows and chimps require biotin for hemoglobin synthesis. Therefore, horses require biotin for hemoglobin synthesis.” The Osherson horse set contains 36 two-premise arguments with the conclusion “horse,” and the mammals set contains 45 three-premise arguments with the conclusion “all mammals.” Participants ranked each set of arguments in increasing strength by sorting cards.

The induction task concerning cities was conducted by Kemp and Tenenbaum (2009). Participants were presented a scenario where Native American artifacts can be found under most large cities, and some kinds of artifacts are found under just one city while others are under a handful of cities. An example inductive argument is: “Artifacts of type X are found under Seattle and Boston. Therefore, artifacts of type X are found under Minneapolis.” There were 28 two-premise arguments with Minneapolis as the conclusion, 28 with Houston as the conclusion, and 30 three-premise arguments with “all large American cities” as the conclusion. These arguments were ranked for strength by sorting cards, in a method intended to mimic the Osherson tasks.

5.4 Bayesian property induction

We used a Bayesian model of property induction that takes a structured representation (sparse graph, tree, space, or raw covariance) and uses it to make predictions in the property induction tasks (Kemp & Tenenbaum, 2009) (see also Heit, 1998). We refer the reader to Kemp and Tenenbaum (2009) for a fuller treatment, but the basics are described here.

Inductive arguments involve questions of the form: Objects X have property f , therefore, how likely is it that objects Y have property f ? The first set of objects X are the premise categories, and the second set Y is the conclusion category. The argument concerns a new feature f which is observed on the set f_X to be equal to the label l_X (here $l_X = [1, 1, 1]$ since all the premise objects have this property). With this notation, the strength of an inductive argument can be modeled as the probability the property is true for the conclusion $f_Y = 1$ given the premise labels $f_X = l_X$, or

$$P(f_Y = 1 | f_X = l_X) = \sum_{f: f_Y=1} P(f | f_X = l_X), \quad (5)$$

where the posterior $P(f | f_X = l_X)$ is the distribution on full instantiations of the binary feature

vector f , given the labeled premises. This distribution can be computed by Bayes' rule

$$P(f|f_X = l_X) = \frac{P(f_X = l_X|f)P(f)}{\sum_f P(f_X = l_X|f)P(f)}. \quad (6)$$

The likelihood $P(f_X = l_X|f) \propto 1$ if the label is consistent with the feature, and otherwise it is 0. The prior distribution $P(f)$ is instantiated by each of the different models. It follows that the strength of the inductive argument can be computed as

$$P(f_Y = 1|f_X = l_X) = \frac{\sum_{f: f_Y=1, f_X=l_X} P(f)}{\sum_{f: f_X=l_X} P(f)}. \quad (7)$$

Intuitively, this is the weighted fraction of features consistent with the premises and conclusion, compared to those consistent with just the premises. Each feature is weighted by its prior probability $P(f)$. Given that each model specifies a Gaussian distribution on features rather than a distribution on binary features, this prior $P(f)$ was approximated by a large number (10^6) of continuous Gaussian samples, thresholded at the mean (0) to create binary features.

5.5 Results

The models were evaluated on their ability to predict argument strength as evaluated by participants. Predictive performance was evaluated as the correlation between people's rankings and the models' evaluation of argument strength. The structural sparsity predictions using $\beta = 8$ for mammals and $\beta = 4$ for cities are highlighted in Fig. 6 and compared with the other models. For mammals, structural sparsity predicted human ratings about as well as the tree form (Fig. 6A-iv). For cities, structural sparsity fit much better than a tree and somewhat worse than a 2D space (Fig. 6B-iv). The raw covariance model fit particularly well for cities, since the similarity data was already based on 2D drawings, but did not fit as well for the mammals data sets.

It is noteworthy that structural sparsity fit well across the range of values of its sparsity parameter (Table 3), and that these different sparsity settings result in qualitatively different structures for the mammals. For low sparsity ($\beta = 1$, Fig. 7D), the model groups similar mammals together, but it is otherwise difficult to interpret. For higher sparsity ($\beta = 8$, Fig. 7A), the structure closely mimics the tree-based structural form (Fig. 7B). Structural sparsity can predict the human

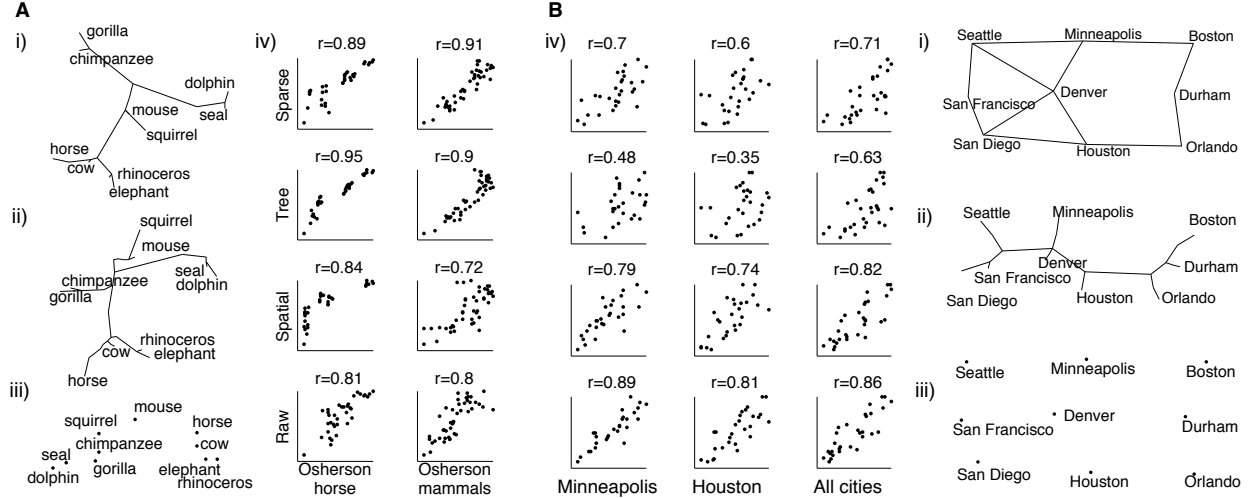


Figure 6: Comparing models and human responses regarding property induction for mammals (A) and cities (B). A dot in (iv) is an inductive argument; for instance, the bottom-left dot for the “Osherson horse” data set is the argument “dolphins and seals have this property, therefore horses do.” Each dot represents a different set of premises, followed by the conclusion category labeled on the x-axis (horse, all mammals, Minneapolis, Houston, or all cities). Argument strength for the models (x-axis) is plotted against mean rank of strength across participants (y-axis), with the correlation coefficient r shown above. Predictions were compared for different types of structures, including those learned with structural sparsity (i), trees (ii), and 2D spaces (iii). The arguments about mammals mention only the 10 mammals shown in (A), although predictions were made by using the full structures learned for 50 mammals. Here only the subtrees (or space) that contains the pairwise paths between these 10 mammals are shown.

judgments using either structure, suggesting that a pristine tree is sufficient, but not necessary, to explain the human judgments.⁷

6 General Discussion

As illustrated across a range of data sets, structural sparsity can discover qualitatively different organizing structures. Rather than predefining a set of forms from which it selects, the model considers a continuum of hypotheses with a simplicity bias towards sparsity. Despite this weaker inductive bias, the model can nevertheless recover distinct structures such as trees, rings, and chains for tightly organized domains. Unlike the structural forms model (Kemp & Tenenbaum, 2008), structural sparsity can learn exceptions to forms which often carry important semantic meaning. If a given domain is more loosely organized, such as artifacts, the model learns a complex organization

⁷Similar results were obtained by learning sparsely graphs without latent cluster variables in Lake and Tenenbaum (2010).


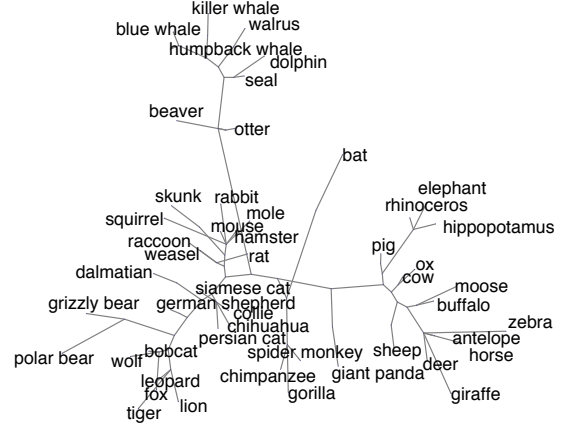

A**B****C****D**

Figure 7: Learned structures to represent the mammals data set. Structural sparsity with $\beta = 8$ (A), tree-based structural form (B), 2D spatial structural form (C), and structural sparsity with $\beta = 1$ (D). Stronger edges are displayed as shorter.

β	Osherson horse	Osherson mammals	Minneapolis	Houston	All cities
1	0.94	0.87	0.70	0.59	0.71
2	0.95	0.86	0.70	0.58	0.71
4	0.93	0.84	0.70	0.61	0.71
6	0.91	0.89	0.69	0.51	0.68
8	0.89	0.91	0.69	0.51	0.68

Table 3: Correlations between how people and the structural sparsity model judge inductive strength, for several tasks concerning mammals and cities. Rows of the table indicate different values of the sparsity parameter β .

that does not fit an easily recognizable form. Once a domain structure is learned, it can also be used to predict the extension of novel properties. Where previous work showed a double dissociation between a tree and a 2D form for matching people’s inferences about animals and cities (Kemp & Tenenbaum, 2009), structural sparsity can learn different types of representations that support different patterns of prediction.

Structural sparsity is related to machine learning approaches for dimensionality reduction and feature prediction that typically operate in more restricted settings. Some algorithms for non-linear dimensionality reduction can be interpreted as using the same graphs and likelihood model as structural sparsity (Lawrence, 2011, 2012). But these algorithms do not have latent clusters and commonly use a fixed graph created by adjoining nearby objects. There are also efficient algorithms for learning sparse Gaussian graphical models, either without latent variables (Banerjee et al., 2008; Lake & Tenenbaum, 2010) or without fully recovering a latent topology (Chandrasekaran, Sanghavi, Parrilo, & Willsky, 2011). Other algorithms learn trees with latent variables (Choi, Chandrasekaran, & Willsky, 2010; Choi, Tan, Anandkumar, & Willsky, 2011) or jointly learn two graphs – one for objects and one for features – although without latent variables (Kalaitzis, Lafferty, & Lawrence, 2013). While these algorithms are computationally efficient and successful in the settings they consider, they do not attempt to explain how people can learn different types of organizing structure for different domains. Structural forms showed how the type of structure can be learned by discrete selection, and here we showed that different forms can arise through a general bias towards sparsity.

Structural sparsity also brings a new perspective to an old debate in cognitive science between symbolic (Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010) versus emergent (McClelland et al., 2010) approaches to knowledge representation. The symbolic tradition uses classic knowledge structures including graphs, grammars, and logic, viewing these representations as the most natural route towards the richness of thought (Griffiths et al., 2010; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). The competing emergent tradition views these structures as epiphenomena: they are approximate characterizations that do not play an active cognitive role. Instead, cognition emerges as the co-occurrent consequence of simpler processes, often operating over vector spaces and distributed representations (McClelland, 2010; McClelland et al., 2010). This debate has been particularly lively with regards to conceptual organization, the domain studied here. The

structural forms model (Kemp & Tenenbaum, 2008) has been criticized by the emergent camp for lacking the necessary flexibility for many real domains, which often stray from pristine forms (McClelland et al., 2010). The importance of flexibility has motivated emergent alternatives, such as a connectionist network that maps animals and relations on the input side to attributes on the output side (Rogers & McClelland, 2004). As this model learns, an implicit tree structure emerges in its distributed representations. But those favoring explicit structure have pointed to difficulties: it becomes hard to incorporate data with direct structural implications like “A dolphin is not a fish although it looks like one” (Kemp & Tenenbaum, 2009), and latent objects in the structure support the acquisition of superordinate classes such as “primate” or “mammal” (Xu & Tenenbaum, 2007). Structural sparsity shows how these seemingly incompatible desiderata could be satisfied within a single approach, and how rich and flexible structure can emerge from a preference for sparsity.

While the sparsity model was formulated at Marr’s computational level (Marr, 1982), any complete understanding must also extend to the algorithmic and implementation levels. Explaining how this model, and other models at the computational level, could be approximated by neural circuits is an important challenge, but some unique features of structural sparsity suggest a way forward. Making an analogy between nodes in the model and a layer within a single neocortical column, the brain seems to have the necessary machinery to produce permissive connectivity in principle, but sparse connectivity in practice. In this microcircuit of cortex, the potential synaptic matrix has been likened to a “tabula rasa,” where any two cells are close enough to form at least one synapse (Kalisman, Silberberg, & Markram, 2005). But the actual synapses are highly sparse, and rather than forming synapses at random, two cells seem to make a discrete connectivity choice by forming multiple synapses together (Fares & Stepanyants, 2009). This analogy is superficial without an understanding of inference, cluster nodes, and whether graph nodes better correspond to single cells or populations (Rumelhart, McClelland, & the PDP Research Group, 1986), but we see this notion of shared structural sparsity between the model and the brain as suggestive and worth exploring.

A final puzzle is to understand how a more explicit notion of structural form might arise. For some domains, there may be no suitable form to characterize the learned outcome of structural sparsity (e.g., Fig. 5B), but for other domains, the learned structure can approximate the pristine output of a graph grammar. As the model currently stands, form is not explicitly identified with

a label (“ring,” “tree,” etc.), although such a label might be necessary for genuine “conceptual change” (Carey, 2009). An explicit structural form may also serve as a basis for analogy between domains, as a hypothesis about causal structure within a domain, or as a vehicle for easier communication of larger domain structure (“left” or “right” in a chain and “clockwise” or “counterclockwise” in a ring). Somehow the mind can grasp structure at this level of intuitive description, without placing the bounds of learning around a small set of forms.

How can both explicit form and flexibility be captured in a single computational approach, and can we extend the work developed here in this direction? One possibility is to extend the approach of Kemp and Tenenbaum (2008) to achieve greater flexibility by allowing exceptions to the grammar-based forms. As an alternative to top-down search, in the spirit of the work presented here, form recognition could proceed from the bottom-up by first learning a sparse representation, and then by analyzing it to see if one or more forms fit sufficiently well. For cases like the color circle where the sparse graph exactly instantiates a form (Fig. 4B), form could be identified by simply checking the grammatical definition of a form (see sketch in Fig. 8). For example, a ring could be identified as any connected graph with two edges per node. In fact, there is behavioral support for a related mechanism: people seem to be capable of identifying the abstract form of sparse relational graphs, when the edges are presented sequentially as social interactions (Kemp, 2007; Kemp, Goodman, & Tenenbaum, 2008). Recognizing forms with exceptions (a tree with cross-branch relations) or forms with more complex structures (“a ring of trees” or “a tree in a ring”, Fig. 1B) may require probabilistic inference over a more sophisticated description language, such as probabilistic predicate logic (Richardson & Domingos, 2006) or probabilistic programs (Goodman, Mansinghka, Roy, Bonawitz, & Tenenbaum, 2008). In any case, if sparsity is as powerful as our results suggest, this additional mechanism may provide a relatively modest step in facilitating the rare yet pivotal conceptual leaps in science and development. A criterion for judging any model of learning is how much richness is gained from initial knowledge to acquired knowledge. By this measure, sparsity could play the primary theoretical role in a larger explanation of qualitative structural change.

Logical law	Interpretation
1. $\forall x \forall y \neg R(x, y)$	No edges between cluster nodes.
2. $\exists x \exists! y R(x, y)$	At least one cluster node has just one edge.
3. $\forall x \forall y R(x, y)$ $\rightarrow \exists! z [z \neq y \wedge R(x, z)]$	Exactly 2 edges per cluster node.
4. $\forall x \forall y R(x, y)$ $\rightarrow (\exists! z [z \neq y \wedge R(x, z)]$ $\vee \neg \exists z [z \neq y \wedge R(x, z)])$	Exactly 1 or 2 edges per cluster node.
5. $\forall x \forall y T(x, y)$	At least one path between cluster nodes.
6. $\forall x \forall y \forall z [R(x, y) \wedge R(x, z)$ $\wedge y \neq z] \rightarrow \neg T_{\setminus x}(y, z)$	There are no cycles.

Figure 8: Sketch of an additional mechanism for identifying structural form by analyzing the output of the structural sparsity model (inspired by Kemp et al., 2008). Forms are defined as conjunctions of laws that operate on the graph output: “Clusters” : Law 1; “Chain” : 2, 4 and 5; “Ring” : 3 and 5; “Tree” : 5 and 6. By applying these rules, the process can identify structural sparsity’s output on the color circle as a ring (Fig. 4B) and on mammals as a tree (Fig. 7A). The relation $R(\cdot, \cdot)$ indicates that two cluster nodes share an edge. Since edges are undirected, there is an implicit law of symmetry. The quantifiers are “for all” (\forall), “there exists” (\exists), and the non-standard “there exists exactly one” ($\exists!$). The predicate $T_{\setminus x}(\cdot, \cdot)$ is the transitive closure of $R(\cdot, \cdot)$ on the set of all objects excluding x .

7 Acknowledgments

We thank Intel Labs for providing the data set of 200 objects and Venkat Chandrasekaran for helpful discussions. This work was supported by a NSF Graduate Research Fellowship to B.M.L; the Center for Brains, Minds, and Machines funded by NSF Science and Technology Center award CCF-1231216; and the Moore-Sloan Data Science Environment at NYU.

8 Appendix

8.1 Generating data from structure

A data set is a matrix $D = \{f^{(1)}, \dots, f^{(m)}\}$ ($D \in \mathbb{R}^{n_x \times m}$), where the rows correspond to objects and the columns are features $f^{(i)}$. A data set is generated by a structure, parameterized by a symmetric matrix $S \in \mathbb{R}^{n_t \times n_t}$. Each row/column in S corresponds to a node in the graph. The set of object nodes is called X and cluster nodes is called Z , where n_x and n_z are their cardinalities, combining for a total of $n_t = n_x + n_z$ nodes in the graph.

A key property of the structural sparsity model is that S is sparse, meaning that most of its

entries are equal to zero ($s_{ij} = 0$). For the remainder, their values are positive ($s_{ij} = s_{ji} > 0$). The sparsity pattern defines the adjacency matrix of the undirected graph, where a non-zero value for s_{ij} means that nodes i and j share an edge. Given that each object node $x \in X$ connects to only one cluster node, then $s_{xz} > 0$ for exactly one $z \in Z$. There are no self-edges, so the diagonal of S is also zero.

Following Zhu, Lafferty, and Ghahramani (2003) and the setup for the structural forms model (Kemp & Tenenbaum, 2008), we introduce the graph Laplacian Δ . Its off-diagonal elements are given by $-S$ and its diagonal elements are

$$\Delta_{ii} = \sum_j s_{ij}.$$

A generative model for features can then be written as

$$\begin{aligned} P(f^{(k)}|S) &\propto \exp\left(-\frac{1}{4}\sum_{i,j} s_{ij}(f_i^{(k)} - f_j^{(k)})^2\right) \\ &= \exp\left(-\frac{1}{2}f^{(k)\top}\Delta f^{(k)}\right). \end{aligned}$$

This equation highlights why the model favors features that are smooth across the graph. Features are probable when connected objects i and j ($s_{ij} > 0$) have a similar value. The stronger the connection s_{ij} (meaning the larger its value), the more important it is for the feature values to match. As pointed out in Zhu et al. (2003), this distribution is not proper, since adding a constant value to the feature vector does not change its probability. Therefore following Zhu et al. (2003), we define the matrix $J = \Delta + \frac{1}{\sigma^2}I$ and use

$$P(f^{(k)}|S, \sigma^2) \propto \exp\left(-\frac{1}{2}f^{(k)\top}Jf^{(k)}\right), \quad (8)$$

which results in a proper density. This distribution is an n_t dimensional Gaussian with zero mean

$$\begin{aligned} P(f^{(k)}|S, \sigma^2) &= N\left(f^{(k)}|0, J^{-1}\right) \\ &= \frac{1}{(2\pi)^{n_t/2}|J|^{-1/2}} \exp\left(-\frac{1}{2}f^{(k)\top}Jf^{(k)}\right). \end{aligned}$$

This generative model for features can also be derived from a maximum entropy formulation, as shown in Lawrence (2011, 2012). This distribution over features is nearly the same as in the structural forms model, except the forms model adds the diagonal term $\frac{1}{\sigma^2}$ only to the observed nodes.

This distribution is also known as a Gaussian Markov Random Field (Koller & Friedman, 2009). The undirected graph S (and equivalently the precision matrix J) instantiates a set of conditional independence relationships between the variables in the graph

$$s_{ij} = 0 \text{ if and only if } \left(f_i^{(k)} \perp f_j^{(k)} | f_{\setminus \{i,j\}}^{(k)} \right).$$

If nodes i and j do not share an edge, their feature values are conditionally independent when the rest of the feature vector is observed (or equivalently, their partial correlation is 0 when controlling for all other nodes). Intuitively, a shared edge is a path of direct dependence.

8.2 Prior distribution on structure

The prior on structures S has a very simple form

$$P(S, \sigma^2) \propto \exp \left(-\beta(\#S) \right),$$

where $\#S$ is the number of edges in a structure. Equivalently, $\#S = \frac{1}{2} \|S\|_0$, where $\|S\|_0$ known as the ℓ_0 -norm, counts the number of non-zero entries in a matrix. The prior has support on the set of all possible graphs where each object has only one connection (which is to a cluster node), there are no empty cluster nodes, each entry has $s_{ij} \geq 0$, and $\sigma^2 > 0$.

This prior is improper because the values in S have no upper bound. Given that this paper aims to just find a single good structure, this issue can be ignored and the prior becomes a penalty $-\beta(\#S)$ in the model's log-score (Eq. 3). If it was desirable (and tractable) to compute a posterior distribution $P(S|D)$, then this improper prior could be extended to form a proper prior. Following the structural forms model, S could be decomposed into its sparsity pattern S_{pat} (which values are non-zeros) and the edge values S_{val} . The forms model places a prior on each entry in S_{val} , and uses

Laplace’s method (MacKay, 2003) to approximate the integral over the parameter values

$$P(D|S_{pat}) = \int P(D|S_{pat}, S_{val}, \sigma^2) P(S_{val}|S_{pat}) P(\sigma^2) dS_{val} d\sigma^2.$$

This strategy is entirely consistent with the structural sparsity model. In fact, approximating the integral would create an additional force for sparsity, known as the Bayesian Occam’s razor (MacKay, 2003). But our current approach was chosen for simplicity, since we found it unnecessary to have two simultaneous forces driving for sparsity. There are some more formal reasons to support this point. As the number of features grows, the Laplace approximation for the integral above asymptotes to a model score known as the Bayesian Information Criterion (BIC) (Koller & Friedman, 2009; Schwarz, 1978). The BIC score maximizes the likelihood while penalizing the number of parameters, just like the penalization arising from our prior on S , although it differs in a scaling factor. Thus the more complex model, which combines these two sources of sparsity, asymptotes to the simpler model with just a larger value of the sparsity parameter β , soaking up both forces that promote sparsity.

8.3 Model implementation

Computing a posterior distribution over structures $P(S|D)$ is very difficult, and thus we aim to find a single good structure S that maximizes the posterior score (Eq. 3). Our approach to this optimization problem consists of two main routines. The outer-routine searches for the best clustering pattern (partition) of objects into cluster nodes, called *cluster search*. The inner-routine searches for the best graph S without changing the clustering pattern, called *connection search*.

8.3.1 Cluster search

The strategy for cluster search is related to standard methods for structure learning in graphical models (Koller & Friedman, 2009), which evaluate small changes to a graph and often greedily select the best change. Given a hypothesis structure and its cluster pattern, cluster search considers a set of local proposals to split or merge cluster nodes. The best scoring move, whether it is a split or a merge, is chosen at each step. In order to assess the quality of each move, the edges must be re-optimized with the connection search routine which is computationally expensive. But

multiple possible moves can be evaluated in parallel. To help mitigate the problem of local optima, search does not terminate immediately when the best local move decreases the score. Instead, the algorithm terminates after the score decreases several times (we used 5). Search also keeps a tabu list of previously visited cluster partitions that it does not revisit. The best evaluated structure S is returned as the solution.

Initialization. Choosing a good initialization can save a lot of computation, and in some cases, lead to better results. Given that cluster search operates over cluster partitions, we use standard clustering methods to propose a set of reasonable candidates for initialization. We use k-means clustering to choose a set of k clusters. Multiple values of k are tried, where the values of k are evenly spaced on a logarithmic scale from 1 to the number of objects n_x . Each candidate is evaluated and scored with connection search. After determining the best value for k in this initial coarse sampling, the algorithm attempts to narrow in on a better value. Picking the closest, previously-tried values of k above and below the current k , k is further optimized in this range by Fibonacci search. Assuming the score is a unimodal function of k within these bounds, Fibonacci search will find the optimum in this range.

Splitting clusters. To split a cluster node, the objects currently assigned to that node must be divided between the two new clusters. Following the general splitting strategy used in the structural forms algorithm (Kemp & Tenenbaum, 2008), the sparsity algorithm chooses two seed objects at random, assigns one to each cluster, and stochastically distributes the remaining objects by picking whichever seed object is closer in feature space. There is also a low probability of choosing the opposite seed object. Rather than evaluating just one split per node, the algorithm tries several (3) randomly chosen seed objects. Each split must then be optimized with connection search and scored by the objective function (Equation 3). For each step during search, the algorithm is limited in the total number of splits it will consider across all of the cluster nodes (30 for most data sets, but 8 for large data sets like the Senate and 200 Objects).

Merging clusters. To merge two cluster nodes, they must combine their attached objects to form a single cluster node. Rather than trying every combination of merges, each cluster node stochastically, but intelligently, picks another to combine with. To help select the merges, the algorithm first calculates the expected value of the latent features for all cluster nodes, and the probability of merging two nodes decreases with the distances in this feature space. The algorithm

also limits the number of merges it evaluates (30 for most data sets, but 8 for the Senate and 200 objects).

Swapping assignments. In addition to splits and merges, the algorithm also tries to swap the cluster assignments of objects (Kemp & Tenenbaum, 2008). These moves do not compete with splitting and merging during each greedy step of the algorithm. Instead, at regular intervals throughout search (we used every 3 moves), each object tries to change its cluster assignment while leaving all others fixed. It tries all possibilities, but it does not re-learn the sparsity pattern for each possible assignment, which is an expensive computation. Instead, it just re-optimizes the existing edge strengths of each candidate proposal S . If a new parent leads to a better score, then S is re-optimized with the full connection search.

8.3.2 Connection search

Connection search is the sub-routine that searches for the best sparse connectivity pattern, given a assignment of objects to cluster nodes. Defining a function $c(S)$ that extracts the cluster assignment, connection search must solve

$$\operatorname{argmax}_{\sigma^2, \{S: c(S)=w\}} \sum_{i=1}^m \log P(f_X^{(i)} | S, \sigma^2) - \beta(\#S), \quad (9)$$

where the current cluster assignment is denoted as w . Given that features are only observed for object nodes, and even those features can be missing, we use the Structural Expectation-Maximization algorithm (Structural EM or SEM) for learning in the presence of missing data (Friedman, 1997).

Structural Expectation-Maximization. SEM reduces the missing data problem to a sequence of simpler structure learning problems with complete data. Rather than maximizing Eq. 9 directly, the structure at iteration $r + 1$ of SEM maximizes

$$S^{[r+1]}, \sigma^{2[r+1]} \leftarrow \operatorname{argmax}_{\sigma^2, \{S: c(S)=w\}} \sum_{i=1}^m E_{Q_i(f_Z^{(i)}, S^{[r]}, \sigma^{2[r]})} \left[\log P(f_X^{(i)}, f_Z^{(i)} | S, \sigma^2) \right] - \beta(\#S), \quad (10)$$

where the expectation is over the conditional distribution of the missing or latent features ($f_Z^{(i)}$)

given the observed features $(f_X^{(i)})$ and the current structure at iteration r ,

$$Q_i(f_Z^{(i)}, S^{[r]}, \sigma^{2[r]}) = P(f_Z^{(i)} | f_X^{(i)}, S^{[r]}, \sigma^{2[r]}).$$

Each iteration of the Structural EM algorithm is guaranteed to improve the original objective, or the marginal probability of the observed features Eq. 9. For the structural sparsity model, each iteration can be decomposed into an E-step (Step 3 below) which computes expected sufficient statistics. This is followed by a Structural M-step that re-optimizes the structure to fit the new sufficient statistics (Step 4). The algorithm is shown below.

- 1: Initialize $S^{[0]}$ and $\sigma^{2[0]}$
- 2: **for** $r = 0, 1, 2, \dots$ until convergence **do**
- 3: $H \leftarrow \frac{1}{m} \sum_{i=1}^m E_{Q_i(f_Z^{(i)}, S^{[r]}, \sigma^{2[r]})} [f^{(i)} f^{(i)T}]$
- 4: $\{S^{[r+1]}, \sigma^{2[r+1]}\} \leftarrow \underset{\sigma^2, \{S: c(S)=w\}}{\operatorname{argmax}} \log |J| - \operatorname{tr}(HJ) - \frac{\beta}{m} \|S\|_0$
- 5: **end for**

The precision matrix J is implicitly a function of S and σ^2 . Step 4 is intractable, as is the case with most structure learning problems, and we describe how we approximately solve it in the next section. Also, we run traditional EM to convergence at the end of each iteration r of Structural EM (see Alternating SEM-EM Friedman, 1997). Traditional EM is the same algorithm as SEM but with $\beta = 0$ and a fixed sparsity pattern for S across iterations.

Optimization with ℓ_1 heuristic. Exactly solving Step 4 is intractable, since all possible sparsity patterns need to be considered and the number of different sparsity patterns for k clusters is $2^{(k^2-k)/2}$. Instead we use a convex relaxation of the optimization which replaces the ℓ_0 norm with the ℓ_1 norm. This heuristic leads to a sparse solution (Banerjee et al., 2008). The ℓ_1 relaxation penalizes a sum of the parameters instead of their cardinality, such that $\|S\|_0 \approx \lambda \sum_{ij} s_{ij}$ for some constant λ . This leads to a convex optimization problem which can be solved exactly (Lake &

Tenenbaum, 2010):

$$\begin{aligned}
& \underset{J, S, \sigma^2}{\operatorname{argmax}} \log |J| - \operatorname{tr}(HJ) - \frac{\beta\lambda}{m} \sum_{ij} s_{ij} \\
& \text{subject to} \\
& J = \operatorname{diag}\left(\sum_j s_{ij}\right) - S + \frac{1}{\sigma^2} I \\
& s_{ij} = 0, \{i, j\} \notin \mathcal{E} \\
& s_{ij} \geq 0, \{i, j\} \in \mathcal{E} \\
& \sigma^2 > 0.
\end{aligned}$$

The set of allowable edges is denoted as \mathcal{E} . Given that J can be defined as an affine mapping from the variables S and σ^2 , we can make the first equality constraint implicit in the objective function and remove J as a variable (Boyd & Vandenberghe, 2004). By setting the missing edges to be zero and removing them as variables, we have a convex optimization problem with just a positivity constraint on the parameters. Efficient methods for solving a related sparse Gaussian problem with ℓ_1 have been developed (Banerjee et al., 2008; Schmidt, Van Den Berg, Friedlander, & Murphy, 2009), although this problem does not have a positivity constraint or the graph Laplacian interpretation. Algorithms for this related problem typically solve the dual optimization problem, although we found that solving the primal problem was more efficient for the variant defined here. To solve it, we used a projected quasi-Newton algorithm and code developed by Schmidt et al. (2009). For a given solution, many values are zero but others are just small. The final sparsity pattern is chosen by thresholding to approximately maximize the complete-data objective in Step 4 of the Structural EM algorithm. The algorithm repeats the optimization and thresholding a total of three times, at $\lambda = \{.5, 1, 2\}$, and the best is picked according to the same objective. For the largest data sets, $\lambda = 1$.

Derivation of Structural EM. More details of the Structural EM derivation are now described. We unpack the SEM objective function in Eq. 10 and derive the algorithm in the section above. Here, the distribution Q_i implicitly depends on the current hypothesis S , although we drop the dependence in the notation.

$$\begin{aligned}
& \underset{\sigma^2, \{S: c(S)=w\}}{\operatorname{argmax}} \log P(S, \sigma^2) + \sum_{i=1}^m E_{Q_i(f_Z^{(i)})} [\log p(f_X^{(i)}, f_Z^{(i)}; S, \sigma^2)] \\
&= \underset{\sigma^2, \{S: c(S)=w\}}{\operatorname{argmax}} \log P(S, \sigma^2) - \frac{m}{2} \log |J^{-1}| - \frac{1}{2} \sum_{i=1}^m E_{Q_i(f_Z^{(i)})} [f^{(i)T} J f^{(i)}] \\
&= \log P(S, \sigma^2) + \frac{m}{2} \log |J| - \frac{1}{2} \sum_{i=1}^m (\sum_{j=1}^n \sum_{k=1}^n E_{Q_i(f_Z^{(i)})} [f_j^{(i)} f_k^{(i)}] J_{jk}) \\
&= \log P(S, \sigma^2) + \frac{m}{2} \log |J| - \frac{1}{2} \sum_{i=1}^m \operatorname{tr}(E_{Q_i(f_Z^{(i)})} [f^{(i)} f^{(i)T}] J) \\
&= \log P(S, \sigma^2) + \frac{m}{2} \log |J| - \frac{m}{2} \operatorname{tr}(H J) \\
&\propto \log |J| - \operatorname{tr}(H J) - \frac{\beta}{m} \|S\|_0
\end{aligned}$$

Above, J is the precision matrix (Eq. 8) and H is the expectation of the empirical covariance,

$$H \leftarrow \frac{1}{m} \sum_{i=1}^m E_{Q_i(f_Z^{(i)})} [f^{(i)} f^{(i)T}].$$

The E-step (Step 3 in the SEM algorithm) uses this formula for computing H . Afterwards, we can approximately solve the optimization problem shown in the last line of the derivation, which forms the M-step (Step 4 in the SEM algorithm).

References

- Banerjee, O., El Ghaoui, L., & D’Aspremont, A. (2008). Model Selection Through Sparse Maximum Likelihood Estimation for Multivariate Gaussian or Binary Data. *Journal of Machine Learning Research*, 9, 485–516.
- Boyd, S., & Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.
- Carey, S. (2009). *The Origin of Concepts*. New York, New York, USA: Oxford University Press.
- Chandrasekaran, V., Sanghavi, S., Parrilo, P. A., & Willsky, A. S. (2011). Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2), 572–596.
- Choi, M. J., Chandrasekaran, V., & Willsky, A. S. (2010). Gaussian Multiresolution Models : Exploiting Sparse Markov and Covariance Structure. *IEEE Transactions on Signal Processing*, 58(3), 1012–1024.
- Choi, M. J., Tan, V., Anandkumar, A., & Willsky, A. S. (2011). Learning Latent Tree Graphical Models. *Journal of Machine Learning Research*, 12, 1771–1812.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407–428.

- Ekman, G. (1954). Dimensions of color vision. *Journal of Psychology: Interdisciplinary and Applied*, 38, 467–474.
- Fares, T., & Stepanyants, A. (2009). Cooperative synapse formation in the neocortex. *Proceedings of the National Academy of Sciences of the United States of America*, 106(38), 16463–8.
- Friedman, N. (1997). Learning Belief Networks in the Presence of Missing Values and Hidden Variables. *Fourteenth International Conference on Machine Learning (ICML)*.
- Goodman, N. D., Mansinghka, V. K., Roy, D. M., Bonawitz, K., & Tenenbaum, J. B. (2008). Church: A language for generative models. *Uncertainty in Artificial Intelligence*.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: exploring representations and inductive biases. *Trends in Cognitive Sciences*, 14(8), 357–64.
- Heit, E. (1998). A Bayesian analysis of some forms of inductive reasoning. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 248–274). Oxford University Press.
- Heit, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin and Review*, 7(4), 569–92.
- Inhelder, B., & Piaget, J. (1964). *The Early Growth of Logic in the Child*. London: Routledge & Kegan Paul.
- Kalaitzis, A., Lafferty, J., & Lawrence, N. D. (2013). The Bigraphical Lasso. In *Proceedings of the 30th International Conference on Machine Learning*.
- Kalisman, N., Silberberg, G., & Markram, H. (2005). The neocortical microcircuit as a tabula rasa. *Proceedings of the National Academy of Sciences of the United States of America*, 102(3), 880–5.
- Kemp, C. (2007). *The acquisition of inductive constraints*. Unpublished doctoral dissertation, MIT.
- Kemp, C., Goodman, N. D., & Tenenbaum, J. B. (2008). Theory Acquisition and the Language of Thought. *Proceedings of the 30th Annual Cognitive Science Society (CogSci)*.
- Kemp, C., & Tenenbaum, J. B. (2008). The discovery of structural form. *Proceedings of the National Academy of Sciences*, 105(31), 10687–92.
- Kemp, C., & Tenenbaum, J. B. (2009). Structured statistical models of inductive reasoning. *Psychological Review*, 116(1), 20–58.
- Koller, D., & Friedman, N. (2009). *Probabilistic Graphical Models*. MIT Press.

- Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.
- Lake, B. M., & Tenenbaum, J. B. (2010). Discovering Structure by Learning Sparse Graphs. *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*.
- Lawrence, N. D. (2011). Spectral Dimensionality Reduction via Maximum Entropy. *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 51–59.
- Lawrence, N. D. (2012). A Unifying Probabilistic Perspective for Spectral Dimensionality Reduction: Insights and New Models. *Journal of Machine Learning Research*, 13, 1609–1638.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444.
- MacKay, D. J. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge, UK: Cambridge University Press.
- Markman, E. M. (1989). *Categorization and Naming in Children*. Cambridge, MA: MIT Press.
- Marr, D. C. (1982). *Vision*. San Francisco, CA: W.H. Freeman and Company.
- McClelland, J. L. (2010). Emergence in Cognitive Science. *Topics in Cognitive Science*, 2(4), 751–770.
- McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences*, 14(8), 348–56.
- Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, 97(2), 185–200.
- Osherson, D. N., Stern, J., Wilkie, O., Stob, M., & Smith, E. E. (1991). Default probability. *Cognitive Science*, 15, 251–269.
- Palatucci, M., Pomerleau, D., Hinton, G., & Mitchell, T. (2009). Zero-Shot Learning with Semantic Output Codes. In Y. Bengio, D. Schuurmans, & J. Lafferty (Eds.), *Advances in Neural Information Processing Systems (NIPS)*.
- Richardson, M., & Domingos, P. (2006). Markov logic networks. *Machine Learning*, 62(1-2), 107–136.
- Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior*, 14(6), 665–681.

- Rivera, M., & Lake, J. (2004). The ring of life provides evidence for a genome fusion origin of eukaryotes. *Nature*, *431*, 152–155.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition: A Parallel Distributed Processing Approach*. Cambridge, MA: MIT Press.
- Rumelhart, D. E., McClelland, J. L., & the PDP Research Group. (1986). *Parallel Distributed Processing: Explorations in the microstructure of cognition. Volume I*. Cambridge, MA: MIT Press.
- Schmidt, M., Van Den Berg, E., Friedlander, M. P., & Murphy, K. (2009). Optimizing Costly Functions with Simple Constraints: A Limited-Memory Projected Quasi-Newton Algorithm. *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics*, *6*(2), 461–464.
- Shepard, R. N. (1980). Multidimensional scaling, tree-fitting, and clustering. *Science*, *210*, 390–398.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to Grow a Mind: Statistics, Structure, and Abstraction. *Science*, *331*(6022), 1279–85.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B*, *58*(1), 267–288.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*(2), 245–272.
- Zhu, X., Lafferty, J., & Ghahramani, Z. (2003). *Semi-supervised learning: From Gaussian fields to Gaussian processes* (Tech. Rep. No. CMU-CS-03-175).